

The GEO glossary

The GEO glossary is a precise definition of every term you'll meet in AI-visibility work, from parametric memory to owned-citation rate. Each entry is self-contained, because sloppy vocabulary is how bad advice spreads — and each one links to itself, so you can point a colleague at a single term.

-
- Written for people new to the subject: no term is defined using another you haven't met yet.
 - The method behind the vocabulary is **the field guide** (mentionbeat.com/guide); the page-level rubric is **the checklist** (mentionbeat.com/checklist).
 - Copy it, edit it, put it in your own wiki — the Markdown version is at the bottom.

GEO (Generative Engine Optimization)

The discipline of making content likely to be surfaced, quoted, and cited by generative AI systems — ChatGPT, Claude, Gemini, Perplexity, Google's AI answers.

AEO (Answer Engine Optimization)

A near-synonym for GEO that emphasizes direct-answer surfaces. In practice the two are used interchangeably.

AI-SEO

An informal umbrella for GEO/AEO plus the classic SEO that underpins it — because live retrieval runs on search indexes, an unindexed page can't be retrieved.

Parametric memory

What a model learned during training — facts that appeared repeatedly and consistently in its training data. Not directly editable, updated only at the next training cycle. **The "slow clock."**

Grounding / RAG

Retrieval-Augmented Generation: at query time the engine retrieves passages from the web and answers from them, usually with citations. Updatable in hours to days. **The "fast clock."**

Grounded vs. parametric answer

An answer built from retrieved sources (usually cited) versus one produced from training memory. They measure different channels and must never be averaged together.

Passage / chunk

The unit of retrieval — a self-contained span of text, not a whole page. The thing that earns a citation is a paragraph that makes sense alone.

Chunkability

How well content survives being split into self-contained retrievable passages: one idea per section, no orphan pronouns, the answer stated first.

Entity

A thing a model reasons about — a product, company, concept, or person — with typed relationships to other things (category, maker, alternatives). LLMs reason over entities, not keyword strings.

Knowledge graph

Structured entity-and-relationship data (Wikidata, Google Knowledge Graph) that feeds both model training and live grounding. The base layer of entity visibility.

sameAs

The schema.org property declaring that several URLs refer to the same entity — the literal machine instruction "these are all me." Cheap, durable, frequently omitted.

Product brief (single source of truth)

The one approved fact sheet a subject's content is generated from — canonical name, category, definition, specs with units, aliases. Contradictions anywhere downstream of it are the #1 cause of AI hedging and hallucination.

Hub-and-spoke

Site architecture where one canonical page (the hub) is supported by focused pages (spokes), each owning one question, linked both ways with descriptive anchors.

Answer-first

Putting the direct answer in the first sentence under a question-shaped heading, then elaborating. The single strongest "cite me" signal a page can send.

Presence / mention rate

The share of test prompts where the subject is named in the answer. Headline mention rate is measured on category prompts that never name the brand — otherwise you're grading your own exam.

Share of voice

Your mentions as a fraction of all options named on category prompts — visibility relative to the rivals engines actually name.

Owned-citation rate

The share of answers citing one of your own URLs — as opposed to answers that describe you correctly but send authority to someone else's page.

Prompt suite

A fixed, versioned set of test prompts re-run on a schedule to measure visibility. Versioned because changing the suite invalidates every historical comparison.

Judge (LLM-as-judge)

A model used to score answers ("was the brand mentioned? recommended?"). A judge has its own error rate, so it requires calibration against human-labeled samples to be trustworthy.

Wilson interval

A confidence interval for proportions that behaves well at extreme rates and small samples — unlike the naive (Wald) interval, which fails exactly where visibility rates live.

Cluster bootstrap

Resampling at the prompt level rather than the response level when computing confidence intervals — because repetitions of the same prompt are correlated, and treating them as independent understates uncertainty.

GEO lint

An automated pre-publish check that scores a page against the GEO ruleset and blocks publishing on hard errors — the checklist as an enforced contract instead of a hope.

AI Overviews / AI Mode

Google's generative answer surfaces. As of mid-2026, AI Mode is the default search experience — one of the specifics that changes fast, so treat any dated claim about it as perishable.

Search agent vs. training crawler

Two different kinds of bot, two different decisions: search/answer agents fetch pages to build the answer a user sees now (blocking them removes you from AI answers); training crawlers collect corpora for future models (blocking them is a legitimate IP choice). Conflating the two is the expensive robots.txt mistake.

llms.txt

A proposed root Markdown file listing your key pages for LLMs. Not broadly adopted as a ranking input — cheap insurance, never a lever. [Generate yours](#) → (mentionbeat.com/tools/llms-txt)

Cannibalization

When a supporting page out-ranks or out-cites the page that should be canonical — the usual result of fragmenting one subject across many thin pages.

Zero-click influence

The buying influence that happens inside the AI answer, with no visit for analytics to count. Most citations never become a click, which is why presence-in-answers — not referral traffic — is the honest instrument.

[Download the PDF](#)[Read the field guide](#)

The method behind the vocabulary: [the GEO Field Guide](#) · why it is its own discipline: [why GEO](#) · how we measure it: [the methodology](#) · everything else we publish: [resources](#)